

REMARKS

Claims 1-20 are pending in this application with claims 1, 2, 6, 7, 14, 16, 18, 19 and 20 being amended by this response to more clearly define the invention. Applicant respectfully submits that no new matter has been added by these amendments. Support for these amendments can be found in the original claims and throughout the specification, in particular on page 13, lines 1-5 and page 12, lines 28-33.

As stated on page 10 of the Office Action, claim 6 would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims. Applicant respectfully submits that claim 6 has been amended in accordance with these suggestions and thus should be allowable.

Objection to Title

The title of the invention is objected to as not being descriptive. Applicant respectfully submits that the amended title, "System for Managing Operational Failure Occurrences in Processing Devices" is clearly indicative of the invention to which the claims are directed. Consequently in view of this amendment to the title, this objection is satisfied and should be withdrawn.

Rejection of claims 1-5 and 7-20 under 35 U.S.C. 102(b)

Claims 1-5 and 7-20 are rejected under 35 U.S.C. 102(b) as being anticipated by Alexander et al. (U.S. Patent No. 6,189,111).

The present claimed invention provides a system used by individual processing devices of a group of networked processing devices, to manage operational failure occurrences in devices of the group. An interface processor maintains transition information identifying a second processing device for taking over execution of tasks of a first processing device in response to an operational failure of the first processing device. The interface processor also dynamically updates the transition information in response to a change in utilization parameters occurring in another processing device of the group. An operation detector detects an operational

failure of the first processing device. A failure controller initiates execution, by the second processing device, of tasks designated to be performed by the first processing device in response to detection of an operational failure of the first processing device. Applicant respectfully submits that Alexander neither discloses nor suggests the features described by the present invention.

Known failover cluster systems have the problem that multiple nodes may fail to the same back-up node causing further failure because of over-burdened computer resources. Additionally, for multiple node clusters, existing methods require substantial configuration effort to manually configure a back-up processing device. Also, existing methods may increase operational burdens, increase the risk of failure and disruption of service, and increase the cost of re-configuration. "A system according to invention principles provides a processing device failure management system addressing the identified problems and deficiencies," (Specification page 1, line19 to page 2, line 15).

Alexander provides a method for "retrieving critical data structures of memory from a failed computing element for reconstruction on a non-failed computing element, allowing such critical data structures to continue with their original function," (Abstract). Alexander merely shows a static system and is concerned with **recovery of memory** following processing device failure. The present claimed invention, on the other hand, is concerned with the **transfer of processing** upon a failure of a processing device.

Applicant respectfully submits that contrary to the assertions in the Office Action, Alexander neither discloses nor suggests "transition information identifying a second processing device for taking over execution of tasks of a first processing device in response to an operational failure of said first processing device and for dynamically updating said transition information in response to a change in utilization parameters occurring in another processing device of said group" as recited in the present claimed invention. Alexander only discloses a cluster membership service (CLMS) that "will notify one of the other nodes that is designated as a recipient of the harvest of the failed node," (Alexander, column 3, lines 25-31 and column 5, lines 35-

53). Alexander does not contemplate taking account of utilization parameters at all in designating a device to assume operation of processes of a failed device. Additionally, the CLMS of Alexander makes the decision of the designated node after the detection of the failure. This consumes time and also requires a timer for active harvesting. The present claimed invention, on the other hand, provides that a designated node is already predetermined before the failure occurs using “transition information” “identifying a second processing device for taking over execution of tasks of a first processing device in response to an operational failure of said first processing device.” In case of a failure, no decision has to be made and no timer has to be started. The designated node can advantageously take over immediately. Additionally, unlike Alexander which merely provides a “static preselection” in column 6, lines 19-26, the preselection choice of the present claimed invention dynamically changes with the approach, based on CPU, memory and other utilization factors. Consequently, as Alexander provides no 35 USC 112 compliant enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 1 under 35 U.S.C. 102(b) is respectfully requested.

Claim 2 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1. Claim 2 is also considered to be patentable because Alexander neither discloses nor suggests, “said utilization parameters identify utilization of resources including at least one of, (a) memory, (b) CPU and (c) input-output communication, used for performing particular computer operation tasks in normal operation,” as recited in the present claimed invention. Rather, column 3, lines 1-6 of Alexander, cited in the Office Action, merely describe that “memory objects are data structures that represent the state of critical information resources” and include “system resources such as shared memory, semaphores, message queues, processes, memory mappings, cluster file tokens, devices and similar data structures.” Applicant respectfully submits that the memory objects described in Alexander are not equivalent to the utilization parameters described by the present claimed invention. In Alexander these memory objects are the information gathered from the failed node in the “resources harvesting” and then reconstructed at the non-failed node to continue the function. In the present claimed invention, on the other hand, changes in any of these utilization parameters causes an update to the transition

information, which in turn identifies “a second processing device for taking over execution of tasks of a first processing device in response to an operational failure of said first processing device.” Consequently, as Alexander provides no 35 USC 112 compliant enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 2 under 35 U.S.C. 102(b) is respectfully requested.

Claim 3 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1. Claim 3 is also considered to be patentable because Alexander neither discloses nor suggests “said transition information comprises a prioritized list of processing devices for assuming execution of tasks of a first processing device in response to an operational failure of said first processing device and said prioritized list is dynamically updated in response to communication from another processing device of said group,” as recited in the present claimed invention. Contrary to the assertion in the Office Action, column 6, lines 19-23 of Alexander merely describes a **static** pre-selection. The present claimed invention, on the other hand, describes that the pre-selection choice dynamically changes with the approach, based on CPU, memory and other utilization factors. Consequently, as Alexander provides no 35 USC 112 compliant enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 3 under 35 U.S.C. 102(b) is respectfully requested.

Claim 4 is dependent on claims 1 and 3 and is considered to be patentable for the reasons given above in connection with claims 1 and 3.

Claim 5 is dependent on claim 3 and is considered to be patentable for the reasons given above in connection with claim 3.

Claim 7 is dependent on claim 3 and is considered to be patentable for the reasons given above in connection with claim 3. Claim 7 is also considered to be patentable because Alexander neither discloses nor suggests, “said prioritized list is dynamically updated to include state information indicating at least one of, (a) a detected change of state of another processing device of said group from available to

unavailable and (b) a detected change of state of another processing device of said group from unavailable to available,” as recited in the present claimed invention. Contrary to the assertions in the Office Action, Alexander in column 8, lines 40-44 and 59-61, merely describes that the cluster management component is only able to detect the failure of the component (e.g. if the “I am alive” message is not received, the conclusion is that one node failed.) In the present claimed invention, on the other hand, other system parameters will also be distributed, e.g. CPU utilization and memory to the cluster management service. There might be situations where the “I am alive” message is sent by faulty local decisions on the failed node, but the system does not perform well if the system memory is too exhausted to provide services. Thus, unlike Alexander, the present claimed invention provides “said prioritized list is dynamically updated to include state information indicating at least one of, (a) a detected change of state of another processing device of said group from available to unavailable and (b) a detected change of state of another processing device of said group from unavailable to available.” Consequently, as Alexander provides no 35 USC 112 compliant enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 7 under 35 U.S.C. 102(b) is respectfully requested.

Claim 8 is dependent on claim 7 and is considered to be patentable for the reasons given above in connection with claim 7. Claim 8 is also considered to be patentable because Alexander neither discloses nor suggests “said interface processor determines a state of a processing device of said group from state information provided by a different processing device of said group,” as recited in the present claimed invention. Contrary to the assertion in the Office Action, and as described above with respect to claim 7, Alexander in column 8, lines 40-44 and 59-61, merely describes that the cluster management component is only able to detect the failure of the component (e.g. if the “I am alive” message is not received, the conclusion is that one node failed.) The state of the node information from other nodes is received and based on that received information, the decision is made if the system is not alive. As stated above with respect to claim 7, if the node with the failure still sends the “I am alive” message, the system will not fail over. There might be situations where the “I am alive” message is sent by faulty local decisions on the failed node, but the system

does not perform well due to reduced availability of memory capacity, for example. Consequently, as Alexander provides no 35 USC 112 compliant enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 8 under 35 U.S.C. 102(b) is respectfully requested.

Claim 9 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1.

Claim 10 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1.

Claim 11 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1.

Claim 12 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1.

Claim 13 is dependent on claim 1 and is considered to be patentable for the reasons given above in connection with claim 1.

Claim 14 provides a system for use by individual processing devices of a group of networked processing devices, for managing operational failure occurrences in devices of the group. An individual processing device includes a repository including transition information identifying a second processing device for taking over execution of tasks designated to be performed by a first processing device in response to an operational failure of the first processing device. The individual processing device further includes an interface processor for maintaining and updating the transition information in response to a change in utilization parameters occurring in another processing device of the group. The individual processing device also includes an operation detector for detecting an operational failure of the first processing device; and a failure controller for initiating execution, by the second processing device, of tasks designated to be performed by the first processing device in response to detection of an operational failure of the first processing device.

Applicant respectfully submits that contrary to the assertions in the Office Action, Alexander neither discloses nor suggests “a repository including transition information identifying a second processing device for taking over execution of tasks designated to be performed by a first processing device in response to an operational failure of said first processing device; an interface processor for maintaining and updating said transition information in response to a change in utilization parameters occurring in another processing device of said group,” as recited in claim 14 of the present invention.

As stated above, Alexander merely shows a static system and is concerned with **recovery of memory** following processing device failure. As described above, Alexander only discloses a cluster membership service (CLMS) that “will notify one of the other nodes that is designated as a recipient of the harvest of the failed node,” (Alexander, column 3, lines 25-31 and column 5, lines 35-53). Alexander does not contemplate taking account of utilization parameters at all in designating a device to assume operation of processes of a failed device. Column 3, lines 25-31 of Alexander, cited in the Office Action, merely describes the CLMS making the decision of the designated node after the detection of failure. This consumes time and also requires a timer for active harvesting. In the present claimed invention, the transition information is information about a predetermined node that should take over if a specific node fails. This approach also works for n+2, n+3, etc. clusters. Additionally, if the load situation on the predetermined node changes, a different node becomes the predetermined failover node for the specific node. In the present claimed invention, a designated node is already predetermined before the failure occurs, (“transition information”) i.e. “identifying a second processing device for taking over execution of tasks of a first processing device in response to an operational failure of said first processing device.” In case of a failure, no decision has to be made and no timer has to be started. The designated node can advantageously take over immediately. Additionally, unlike Alexander which merely provides a “static preselection” in column 6, lines 19-26, the preselection choice of the present claimed invention dynamically changes with the approach, based on CPU, memory and other utilization factors. Consequently, as Alexander provides no 35 USC 112 compliant

enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 14 under 35 U.S.C. 102(b) is respectfully requested.

Claim 15 is dependent on claim 14 and is patentable for the reasons given above in connection with claim 14. Claim 15 is also patentable because Alexander neither discloses nor suggests “said interface processor communicates with other processing devices of said group to maintain consistent transition information in said individual processing device transition information repositories,” as recited in the present claimed invention. Rather, column 3, lines 1-6 of Alexander, cited in the Office Action, merely describe that “memory objects are data structures that represent the state of critical information resources” and include “system resources such as shared memory, semaphores, message queues, processes, memory mappings, cluster file tokens, devices and similar data structures.” Applicant respectfully submits that the memory objects described in Alexander are not equivalent to the transition information described by the present claimed invention. In Alexander these memory objects are the information gathered from the failed node in the “resources harvesting” and then reconstructed at the non-failed node to continue the function. In the present claimed invention, on the other hand, changes in the transition information causes an identification of “a second processing device for taking over execution of tasks of a first processing device in response to an operational failure of said first processing device.” Consequently, as Alexander provides no 35 USC 112 compliant enabling disclosure of each feature of the present claimed method, withdrawal of the rejection of claim 15 under 35 U.S.C. 102(b) is respectfully requested.

Claim 16 describes a system for use by individual processing devices of a group of networked processing devices, for managing operational failure occurrences in devices of said group. Claim 16 includes features similar to those described above with respect to claim 1 and 14. Therefore, claim 16 is considered to be patentable for the reasons given above with respect to claims 1 and 14.

Claim 17 is dependent on claim 16 and is considered to be patentable for the reasons given above in connection with claim 16. Claim 17 also includes features

similar to those described above with respect to claim 3 and is therefore considered to be patentable for reasons given above in connection with claim 3.

Claim 18 describes a method for use by individual processing devices of a group of networked processing devices, for managing operational failure occurrences in devices of said group. Claim 18 includes features similar to those described above with respect to claims 1 and 2. Therefore claim 18 is considered to be patentable for the reasons given above in connection with claims 1 and 2.

Claim 19 describes a method for use by individual processing devices of a group of networked processing devices, for managing operational failure occurrences in devices of said group. Claim 19 includes features similar to those described above with respect to claims 1 and 2. Therefore claim 19 is considered to be patentable for the reasons given above in connection with claims 1 and 2.

Claim 20 describes a method for use by individual processing devices of a group of networked processing devices, for managing operational failure occurrences in devices of said group. Claim 20 includes features similar to those described above with respect to claim 1 and 14. Therefore, claim 20 is considered to be patentable for the reasons given above with respect to claims 1 and 14.

Claim 6 has been indicated as allowable if rewritten in independent form including the limitations of the base claim and any intervening claims. Claim 6 has been amended to include the features of claims 1 and 5 and this is now in condition for allowance.

In view of the above remarks and amendments to the claims, it is respectfully submitted that there is no 35 USC 112 enabling disclosure in Alexander et al. that makes the present claimed invention unpatentable. Thus, it is respectfully submitted that Alexander et al. does not disclose each claimed feature and does not anticipate the present invention as claimed in claims 1, 14, 16, 18, 19 and 20. As claims 2-5 and 7-13 are dependent on independent claim 1, claim 15 is dependent on claim 14 and claim 17 is dependent on claim 16, it is respectfully submitted that claims 2-5, 7-13,

15 and 17 are similarly not anticipated by Alexander. Therefore, it is respectfully requested that the rejection of claims 1-20 under 35 USC 102(e) be withdrawn.

Having fully addressed the Examiner's rejections, it is believed that, in view of the preceding amendments and remarks, this application stands in condition for allowance. Accordingly then, reconsideration and allowance are respectfully solicited. If, however, the Examiner is of the opinion that such action cannot be taken, the Examiner is invited to contact the applicant's attorney at the phone number below, so that a mutually convenient date and time for a telephonic interview may be scheduled.

Respectfully submitted,
Arnold Monitzer



Date: February 12, 2007

Alexander J. Burke
Reg. No. 40,425

Siemens Corporation
Customer No. 28524
Tel (732) 321 3023
Fax (732) 321 3030
February 9, 2007